

CLAIMS

5

1. A method of connecting a client application to a server 'server' on a cluster 'cluster' having a plurality of servers instantiated thereon, the method comprising:

receiving 'cluster' and 'server' from the client application;

10

sending a first request message to 'cluster' requesting first connection information for connecting to 'server';

receiving from 'cluster' a first reply message containing the requested first connection information;

15

connecting the client application to 'server' on 'cluster' based on the received first connection information, wherein once connected, the client application and 'server' may transact business;

20

determining that the connection to 'server' has failed, 'server' presumably having been moved from a first server of the cluster to a second server of the cluster, the received first connection information corresponding to the first server;

sending a second request message to 'cluster' requesting second connection information for connecting to 'server', the requested second connection information corresponding to the second server;

25

receiving from 'cluster' a second reply message containing the requested second connection information; and

connecting the client application to 'server' on 'cluster' based on the received second connection information, wherein once again connected, the client application and 'server' may again transact business.

2. The method of claim 1 comprising sending each request message as a UDP (user datagram protocol) packet addressed to a UDP address of 'cluster'.

5 3. The method of claim 1 comprising receiving each requested connection information including an address of 'cluster' on which 'server' is listening.

10 4. The method of claim 1 comprising receiving each requested connection information including an address of 'cluster' on which 'server' is listening for packets formatted according to a VIA (Virtual Interface Architecture) protocol.

15 5. The method of claim 1 comprising receiving each requested connection information including each address of 'cluster' on which 'server' is listening and a corresponding protocol associated with the address by 'server'.

20 6. The method of claim 1 comprising connecting the client application to 'server' on 'cluster' at an address of 'cluster' on which 'server' is listening for packets formatted according to a VIA (Virtual Interface Architecture) protocol.

25 7. The method of claim 1 further comprising:
caching the received second connection information in a
cache;
subsequently again receiving 'cluster' and 'server' from the
client application;
retrieving the cached connection information from the cache;
connecting the client application to 'server' on 'cluster' based
30 on the retrieved cached connection information.

8. The method of claim 1 further comprising:
determining whether the connected-to server is in fact 'server'
and if not:

5 sending a new request message to 'cluster' requesting new
connection information for connecting to 'server';
receiving from 'cluster' a new reply message containing the
requested new connection information; and
connecting the client application to 'server' on 'cluster' based
10 on the received new connection information.

9. A computer-readable medium having stored thereon
computer-executable instructions implementing a method of connecting a client
application to a server 'server' on a cluster 'cluster' having a plurality of servers
15 instantiated thereon, the method comprising:
receiving 'cluster' and 'server' from the client application;
sending a first request message to 'cluster' requesting first
connection information for connecting to 'server';
receiving from 'cluster' a first reply message containing the
20 requested first connection information;
connecting the client application to 'server' on 'cluster' based
on the received first connection information, wherein once connected, the client
application and 'server' may transact business;
determining that the connection to 'server' has failed, 'server'
25 presumably having been moved from a first server of the cluster to a second
server of the cluster, the received first connection information corresponding to
the first server;
sending a second request message to 'cluster' requesting
second connection information for connecting to 'server', the requested second
30 connection information corresponding to the second server;

receiving from 'cluster' a second reply message containing the requested second connection information; and

connecting the client application to 'server' on 'cluster' based on the received second connection information, wherein once again connected,
5 the client application and 'server' may again transact business.

10. The medium of claim 9 wherein the method comprises sending each request message as a UDP (user datagram protocol) packet addressed to a UDP address of 'cluster'.

11. The medium of claim 9 wherein the method comprises receiving each requested connection information including an address of 'cluster' on which 'server' is listening.

12. The medium of claim 9 wherein the method comprises receiving each requested connection information including an address of 'cluster' on which 'server' is listening for packets formatted according to a VIA (Virtual Interface Architecture) protocol.

13. The medium of claim 9 wherein the method comprises receiving each requested connection information including each address of 'cluster' on which 'server' is listening and a corresponding protocol associated with the address by 'server'.

14. The medium of claim 9 wherein the method comprises connecting the client application to 'server' on 'cluster' at an address of 'cluster' on which 'server' is listening for packets formatted according to a VIA (Virtual Interface Architecture) protocol.

15. The medium of claim 9 wherein the method further comprises:

 caching the received second connection information in a cache;

5 subsequently again receiving 'cluster' and 'server' from the client application;

 retrieving the cached connection information from the cache;
 connecting the client application to 'server' on 'cluster' based on the retrieved cached connection information.

10

16. The medium of claim 9 wherein the method further comprises:

 determining whether the connected-to server is in fact 'server' and if not:

15

 sending a new request message to 'cluster' requesting new connection information for connecting to 'server';

 receiving from 'cluster' a new reply message containing the requested new connection information; and

20

 connecting the client application to 'server' on 'cluster' based on the received new connection information.